

# Inferring social structure from continuous-time interaction data

Wesley Lee\*  
University of Washington

Bailey K. Fosdick  
Colorado State University

Tyler H. McCormick  
University of Washington

## Abstract

Relational event data, which consist of events involving pairs of actors over time, are now commonly available at the finest of temporal resolutions. Existing continuous-time methods for modeling such data are based on point processes and directly model interaction “contagion,” whereby one interaction increases the propensity of future interactions among actors, often as dictated by some latent variable structure. In this article, we present an alternative approach to using temporal-relational point process models for continuous-time event data. We characterize interactions between a pair of actors as either spurious or that resulting from an underlying, persistent connection in a latent social network. We argue that consistent deviations from expected behavior, rather than solely high frequency counts, are crucial for identifying well-established underlying social relationships. To explore these latent network structures in two contexts: one comprising of college students and another involving barn swallows.

**Keywords:** continuous time, latent network, point process, relational event data

## 1 Introduction

Data consisting of interactions between pairs of actors (i.e. dyads) over time are commonly referred to as relational event data (Butts, 2008) and are being collected at unprecedented rates. As mobile technology becomes nearly ubiquitous, it is increasingly possible to measure social interactions with arbitrarily granular temporal precision. Call detail records (CRDs), for example, provide detailed descriptions of mobile phone interactions on a national scale (Blumenstock, 2012). The widespread adoption of social media has led to a wealth of continuous-time activity data (Sadilek et al., 2012). In addition, in ecology, advances in tracking devices have resulted in an increase in animal telemetry logs documenting movement and interactions between animals (Rutz and Hays, 2009).

---

\*Contact email: wtlee@uw.edu


Relational event data are distinct from traditional, survey-based network measures in two key ways. First, relational events do not directly quantify relationships between actors, but serve as indicators of these relations and embed underlying social structure and dynamics. Second, relational event data are typically obtained from administrative logs or other sources not explicitly designed for research. This feature presents additional statistical challenges, but also means that collecting relational event data is often relatively inexpensive compared to collecting survey data and can allow researchers to study larger populations of actors at higher temporal resolutions and for more extended periods (Watts, 2007).

Analyses of relational event data typically involve representing interactions in data structures congruent with existing discrete-time dynamic social network models (Sarkar and Moore, 2005; Durante and Dunson, 2014; Krivitsky and Handcock, 2014; Sewell and Chen, 2015). To mirror the form of traditional survey network data, interactions are commonly aggregated into a series of weighted adjacency matrices, also called sociomatrices, whose elements represent the number of interactions between each pair of actors within fixed time intervals. Taken together, the sequence of adjacency matrices can be viewed as a weighted network which evolves in discrete-time. Inference on discrete-time dynamic networks is typically performed using models that characterize the manner in which dyadic interactions change from one time interval to the next, often assuming a Markov process whereby the network at time  $t$  only depends on its past states through the network at time  $t - 1$  (Snijders, 2005; Sewell and Chen, 2016).

One issue with the aforementioned aggregation approach is the implicit assumption that unobserved meaningful social (network) relations are easily ascertained from the observed noisy interaction data. The aggregation methods are designed for scenarios in which the data contain complete network information, consisting of the strength of each dyad relation. However, continuous-time interaction data consists only of measurements *when individuals interact*, and the absence of interaction should not be taken as an explicit declaration of no relationship. More generally, interaction counts between dyads should not necessarily be taken as a direct measure of the strength of the underlying dyad relationship. Consider call records in which actor A may call actor B multiple times per week and call actor C on Sunday morning once a month. From our perspective, both of these call patterns may indicate strong relations between actor A and the others. Actor B may, for example, represent a roommate, whereas actor C may represent actor A’s grandparent. Although communication with the grandparent is relatively infrequent compared to the communication with the roommate, the relationship between the actor A and his/her grandparent should be classified as strong and significant.

Another key drawback of performing discrete-time data aggregation when modeling

continuous-time interaction data is that the time intervals are often chosen arbitrarily and conclusions can be greatly impacted by these choices (Sulo et al., 2010; Sekara et al., 2016). In addition, the temporal dynamics of the interactions are possibly lost in the aggregation process depending on the length of the time interval. Consider Figure 1, for example, where proximity data from a barn swallow study has been aggregated over each day (see Sections 1.1 and 5.1 for further details on this data). By aggregating to the level of the day



**Figure 1:** Daily snapshots of aggregated barn swallow interactions. Nodes are colored by sex (females red; males blue) and edges widths are proportional to the number of interactions between each pair. Birds deemed to be unmonitored on a given day (see Appendix B.1) are shown in white.

we preclude the possibility of capturing intricate within-day social dynamics. While smaller time intervals are able to capture more of these dynamics, models with Markov structure may be inappropriate for modeling the resulting series of weighted networks. Conditional dependence among behavior across time would then be more likely to span multiple intervals. Returning to the example of call records, an assumption of Markovian dynamics for summaries of behavior at, for example, an hourly level would be unsuitable since associated individuals are unlikely to call each other on such a frequent basis.

In this article, we propose a continuous-time approach to modeling relational event data that explicitly separates interactions from underlying social relations. Specifically, we conceptualize continuous-time interaction data as representing *manifestations of latent network structure*. The model we propose possesses two distinctive properties. First, rather than viewing the data as direct observations of the network of interest, we assume observed interactions arise from a point processes with propensity influenced by the latent network structure. It is the dynamics of this *underlying* social network that we argue is typically most informative about population social structure and of direct interest to researchers. Second, we avoid decisions on temporal resolution by modeling both the observed interactions and dynamics of the latent network structure in continuous-time. The statistical challenge, therefore, is to infer the underlying network structure and its evolution through time. In this paper we assume interactions and connection in the latent network are undirected, although our methods could readily be extended to handle directed interactions and networks.

Existing work in continuous-time modeling for relational data focuses on *excitatory* point processes (Simma and Jordan, 2010; Blundell et al., 2012; Perry and Wolfe, 2013). Hawkes processes (Hawkes, 1971) are one such process; in a Hawkes process, the occurrence of an event causes an immediate increase in intensity for that process and/or other Hawkes processes. These processes are favored for their ability to model reciprocity and transitivity in relational event data. Despite being used to model continuous-time data, existing point process models assume a static underlying network (Blundell et al., 2012; Linderman and Adams, 2014) or eschew modeling a network altogether, focusing on just modeling the interaction data (DuBois and Smyth, 2010; Simma and Jordan, 2010; Perry and Wolfe, 2013).

In this article, our primary focus is on inferring an underlying dynamic network, as opposed to inferring casual dependencies between the noisy interactions. In this sense, our work parallels that of Scharf et al. (2015) and Linderman and Adams (2014), which aim to infer network relations from individual-level movement and behavior. We posit that the patterns in relational events are primarily governed by temporal and covariate-dependent factors and only secondarily by excitatory triggers. Thus, we model relational events as arising from inhomogeneous Poisson processes with intensities dependent on time, covariates, and an underlying dynamic latent network. In the interest of parsimony, if two individuals lack a relationship in the underlying network, we consider their interactions as the result of spurious behavior and thus use a relatively simple model. Model complexity is instead focused on interactions between pairs of individuals inferred to have a connection in the underlying network.

The remainder of the paper is organized as follows. In Section 1.1, we briefly introduce two applications for our relational event framework involving face-to-face interactions among college students and animal telemetry data. In Section 2, we describe the general framework in detail, and in Section 3 we propose a Bayesian inference procedure for models from this framework. We apply our framework to the two aforementioned datasets in Sections 4 and 5.

## 1.1 Proximity interactions: students and swallows

Using our method, we explore network structure in two settings with data collected using proximity sensors: one involving interactions between college students and another consisting of contacts among barn swallows (*Hirundo rustica erythrogaster*).

First, we consider proximity information for 57 undergraduate students living in a dormitory at MIT during the 2008-2009 school year (Madan et al., 2012). Participants were given Bluetooth-enabled smartphones, and time stamps of Bluetooth (proximity) pings between

nearby phones were recorded. In addition, students were surveyed five times throughout the school year and asked about assorted measures of their health and relationships with other subjects in the study. This dataset has been used to study the spread of infection (Dong et al., 2012) and the proliferation of health-related choices among these undergraduates (Madan et al., 2010). In both studies, researchers found physical contact information derived from the proximity data played a critical explanatory role beyond that of the self-reported relationships. These findings imply that information obtained from human behavior can inherently differ from that in reported relationships, and the relative importance and validity of each information type may vary across applications.

Next, we examine interactions within a population of Colorado barn swallows collected over three days during mating season (Levin et al., 2015, 2016). Barn swallows are small birds often found nesting in man-made structures. Due to their small size and speed in flight, recording proximity-based social interactions by visual observation is infeasible. Proximity loggers have emerged as promising alternative tools for collecting animal interaction data in species of all sizes. The deployment of these tools presents new opportunities to gain understanding of animal social structure. We investigate a set of close encounters between barn swallows outfitted with proximity loggers.

## 2 A Continuous-Time Interaction Framework

We propose a hierarchical modeling framework for continuous-time relational event data in which dyad interactions are modeled using an inhomogeneous Poisson process with intensity dependent on dyadic covariates and the state of the dyad in an underlying latent network. The latent network relation for each dyad is binary and modeled with a slow-changing continuous-time Markov chain. This construction allows our framework to express two properties of primary interest: 1) We avoid directly equating interaction frequency with relational strength in the network by detaching the process of modeling the relational events from the process of modeling the social network. Consequently, actors that are connected in the latent network can assume interaction patterns with varying frequencies. 2) Both the interactions and the network are modeled directly in continuous-time so that no choices about temporal resolution of data aggregation are required. Furthermore, we retain both the flexibility to capture fine temporal dynamics and the computational benefits of modeling the network evolution as Markovian.

For every pair of actors  $(i, j)$ , denote the vector of interaction event times as  $\mathbf{t}^{ij} \equiv \{t_1^{ij}, \dots, t_{N_{ij}}^{ij}\}$  and the time interval on which these events have the potential to occur by  $\mathbf{T}^{ij} \equiv [T_1^{ij}, T_2^{ij})$ . Additionally, denote the set of dyadic, possibly time-varying, covariates

as  $x_{ij}(t)$ . For the MIT data, covariates of potential interest include whether students reside on the same floor of the dormitory or are in the same year of school. For the swallow data, covariates include sex pairings and similarities between phenotypic traits.

We model the vector of interaction times  $\mathbf{t}^{ij}$  for dyad  $(i, j)$  using an inhomogeneous Poisson process with intensity  $\lambda_{ij}(t)$ , representing the instantaneous rate of an interaction occurring between actors  $i$  and  $j$  at time  $t$ . The log probability of events occurring at times  $\mathbf{t}^{ij}$  for dyad  $(i, j)$  is then given by

$$\log p(\mathbf{t}^{ij} | \lambda_{ij}(t)) = \sum_{t \in \mathbf{t}^{ij}} \log \lambda_{ij}(t) - \int_{\mathbf{T}^{ij}} \lambda_{ij}(t) dt. \quad (1)$$

Conditional on the underlying intensity function, the presence/absence of interactions in disjoint time intervals are independent. Additionally, we assume interactions for different dyads arise independently conditional on their intensity functions. In contrast to Hawkes processes, dependence between dyad interactions is therefore entirely captured by the dyad intensities. Using  $y_{ij}(t)$  as an indicator of whether a network connection exists between  $i$  and  $j$  at time  $t$ , we model each intensity  $\lambda_{ij}(t)$  as dependent on the time  $t$ , dyad-specific covariates  $x_{ij}(t)$ , and the latent connection  $y_{ij}(t)$ .

Introducing time and covariate dependencies of the event data into the Poisson intensity function modulates the influence of these factors in the estimated network and allows the model to differentially weigh interactions as evidence of network structure. Interactions occurring at specific times and/or under certain circumstances may be much more probable in the presence of a network connection rather than in the absence thereof, and should be accounted for accordingly. We structure  $\lambda_{ij}(t)$  as follows:

$$\lambda_{ij}(t) = \left( 1 + m(t, x_{ij}(t)) y_{ij}(t) \right) \times \left[ w(t, x_{ij}(t)) + a(t, x_{ij}(t)) y_{ij}(t) \right], \quad (2)$$

and denote the corresponding sets of model parameters for functions  $\{m, w, a\}$  as  $\{c_m, c_w, c_a\}$ . Assuming no connection in the latent network,  $y_{ij}(t) = 0$ , the intensity function reduces to a baseline intensity  $w(t, x_{ij}(t))$ , dependent on time and dyad-specific covariates. The function  $w$  corresponds to the rate of “spurious” interactions, which result when actors interact due to circumstance. For example, students in the MIT dormitory may be in close proximity to one another at night due solely to room assignments. We leave the exact form of  $w$ , as well as the other functions, to be specified on an case-by-case basis using domain expertise and descriptive summaries of the data.

Behavior for connected dyads can differ from the baseline in two ways: they may generally

tend to interact more often, through the multiplicative factor  $m$ , and their interactions may have a different pattern, through an additive adjustment to the baseline  $a$ . When inferring the latent network, interactions in line with expected behavior for connected dyads but not expected behavior for unconnected dyads are suggestive of a latent connection. Thus, both evidence of generally higher activity levels than baseline and interactions that deviate from the baseline pattern according to our adjustment  $a$  suggest an underlying relationship.

Poisson processes model interactions as instantaneous events in time. However, for many potential applications, including both the MIT and barn swallow proximity data, encounters can last for non-negligible durations. In these cases, we remove the time while the event is taking place from the time interval  $\mathbf{T}^{ij}$  for each given pair, preventing the model from modeling the probability of another event while an interaction is taking place and assuring that the intensity is calibrated to the appropriate time window. Future work could extend our modeling framework to use the duration length as a potential source of information about future interactions and/or network relations.

The latent network  $\mathbf{Y} \equiv \{y_{ij}(t) : i, j \in \{1, \dots, n\}, i \neq j, t \in \mathbf{T}^{ij}\}$  can be decomposed into dyad-specific paths  $\{y_{ij}(t) : t \in \mathbf{T}^{ij}\}$  representing the state of a connection between dyad  $(i, j)$  over time. These paths  $\{y_{ij}(t)\}$  are modeled as independent stationary continuous-time Markov chains (CTMCs) taking values in  $\{0, 1\}$ . Each path  $y_{ij}(t)$  is characterized by a sparsity parameter  $s > 0$  and a transition parameter  $q > 0$ . The probability of transitioning from one network state to another in any interval of length  $t$  is given by the entries of the following transition matrix:

$$P = \begin{matrix} & \begin{matrix} 0 & 1 \end{matrix} \\ \begin{matrix} 0 \\ 1 \end{matrix} & \begin{bmatrix} (1-s) + se^{-qt} & s - se^{-qt} \\ (1-s) - (1-s)e^{-qt} & s + (1-s)e^{-qt} \end{bmatrix} \end{matrix}. \quad (3)$$

This Markov chain has stationary distribution equal to  $(1-s, s)$ , meaning that, asymptotically,  $y_{ij}(t) = 1$  proportion  $s$  of the time. Thus, the parameter  $s$  governs the overall propensity for a connection between any dyad, while  $q$  governs the speed at which the network changes, i.e. relations form/dissolve. Modeling each binary relation as a CTMC with small  $q$  values yields a slow-changing network.

In addition, we allow the CTMC parameters to depend on non-time-varying dyadic covariates, i.e.  $s = s(x_{ij})$  and  $q = q(x_{ij})$  with respective sets of parameters  $\{c_s, c_q\}$ . Allowing for the sparsity parameter to depend on dyadic covariates allows for the introduction of some network structure, such as block models. For example, one could cluster individuals in the inferred network into known communities (such as floors in the MIT data) by allowing

the sparsity parameter to positively dependent on whether or not individuals in each dyad shared the same community.

The hierarchical nature of the model allows us to separate the process of modeling a network from the process of modeling interactions. The interaction model acts as a low-pass filter on the data, distinguishing between spurious interactions and those suggestive of a meaningful relationship. Evidence of network relations between actors is more generally interpreted through deviance from the expected baseline level of activity rather than only through high raw interaction counts. The hidden CTMC under the interaction model smooths evidence across time so that relationships are thus characterized by consistent, prolonged deviations from baseline behavior.

### 3 Bayesian Inference

We propose a Bayesian sampling procedure for estimating the parameters of the model  $\boldsymbol{\theta} \equiv \{c_m, c_w, c_a, c_s, c_q\}$  and the latent network  $\mathbf{Y} = \{y_{ij}(t) : i, j \in \{1, \dots, n\}, i \neq j, t \in \mathbf{T}^{ij}\}$ . After appropriately eliciting priors, we are interested in the posterior distribution of the parameters and network conditional on the data  $D \equiv \{\mathbf{t}^{ij} : i, j \in \{1, \dots, n\}\}$ . Although the posterior distribution is analytically intractable, we can approximate it to an arbitrary degree of accuracy by taking samples from the posterior with a Markov chain Monte Carlo (MCMC) algorithm. MCMC algorithms involve constructing a Markov chain with stationary distribution equal to the posterior distribution of the parameters and network given the data.

Inference proceeds by first sampling the parameters and then estimating the network. The posterior distribution of the parameters  $\boldsymbol{\theta}$  given the data  $D$  is defined

$$p(\boldsymbol{\theta}|D) \propto \left[ \prod_{(i,j)} p(\mathbf{t}^{ij}|\boldsymbol{\theta}) \right] \times p(\boldsymbol{\theta}). \quad (4)$$

After obtaining a sample of size  $K$  from this posterior, the posterior probability that dyad  $(i, j)$  is connected in the network at time  $t$  can then be approximated as


$$p(y_{ij}(t) = 1|D) \approx \frac{1}{K} \sum_{k=1}^K p(y_{ij}(t) = 1|D, \boldsymbol{\theta}^{(k)}), \quad (5)$$

where  $\boldsymbol{\theta}^{(k)}$  is the  $k$ th sample of  $\boldsymbol{\theta}$  from the posterior. The main challenge in this procedure is calculating  $p(\mathbf{t}^{ij}|\boldsymbol{\theta})$ , the probability of a dyad's interactions conditional only on the parameters, since doing so requires marginalizing over the the dyad's network path  $\{y_{ij}(t) : t \in \mathbf{T}^{ij}\}$ . This marginalization is intractable for our model due to the continuous-time nature of the



Poisson process data and the infinite number of possible network paths.

To circumvent this integration issue, we approximate the probability of each dyad's interactions given the network and parameters by discretizing the underlying network path and restricting potential transitions to a fixed number of times dependent on the observed interactions. Namely, we partition the full time interval  $\mathbf{T}^{ij}$  based on the data, such that each (sub)interval consists of a single interaction and the surrounding times that are closer to this interaction than any other (see Figure 2). We then restrict the network relation to be constant within each of these intervals for the purpose of calculating the probability of the observed interactions. The state of the network edge under this partition is characterized by Markov behavior from the midpoint of one subinterval to the next. The object of inference then changes from a continuous-path  $\{y_{ij}(t) : t \in \mathbf{T}^{ij}\}$  to the approximation  $\{y_{ij}(t_l^{ij,*})\}_{l=1}^{N_{ij}}$ , where  $t_l^{ij,*}$  is the midpoint of the interval containing  $t_l^{ij}$ . To infer the state of the path at



**Figure 2:** Sample partitioning of  $(T_1, T_2)$  based on four interactions. Above the horizontal time axis are the time points  $\{t_1, t_2, t_3, t_4\}$  at which events were observed. Below the axis are the midpoints of each interval  $\{t_1^*, t_2^*, t_3^*, t_4^*\}$ . In the interest of parsimony, the network is estimated at each of the midpoint locations and assumed constant within each interval when calculating the probability of the observed data.

other times  $t$ , we simply interpolate between the probabilities of a relation from the nearest two midpoints. Since only one interaction is recorded within each interval and the network is assumed to be slow-changing, we do not expect the underlying network to change repeatedly between interval midpoints.

The procedure we use is not the only possible discretization procedure; for example, one might partition all time intervals  $\mathbf{T}^{ij}$  into preset, equally spaced intervals. However, one advantage of our chosen method is that it automatically trades off between parsimony and flexibility based on the interaction behavior for a dyad. Tying potential state changes for a network path to interactions focuses modeling effort on periods where there is the most information (arising through interactions) in the data and spends less effort modeling periods of low activity, where there is comparatively little information in the data.

We can view  $\{y_{ij}(t_l^{ij,*})\}_{l=1}^{N_{ij}}$  as a hidden Markov model (HMM) with “emissions”  $t^{ij}$  arising from a Poisson process with intensity  $\lambda_{ij}(t)$ . That is, we see evidence of the evolution

of the approximated network path, which itself is unobserved, through the relational event data observed in each interval, and borrow strength across intervals to ensure the path is appropriately slow-changing. The “forward-backward algorithm” is a method for estimating the hidden states in an HMM (Rabiner, 1989). Using the forward variables from this algorithm, we can marginalize over the latent path to obtain an approximation to the probability  $p(\mathbf{t}^{ij}|\boldsymbol{\theta})$ . We can then approximate the posterior  $p(\boldsymbol{\theta}|D)$  in (4), and, in turn, obtain posterior samples  $\{\boldsymbol{\theta}^{(k)}\}$  and subsequently estimate  $\{y_{ij}(t) : t \in \mathbf{T}^{ij}\}$  at the midpoints of the intervals  $\{t_l^{ij,*}\}_{l=1}^{N_{ij}}$  using the forward-backward algorithm.

We use Stan (Carpenter et al., 2016) to sample from the posterior distribution of the parameters and latent network. See the Appendices A.2 and B.2 for more details on this procedure.


## 4 Proximity Interactions among College Students

In this section, we analyze interactions between college students from the MIT Social Evolution dataset (Madan et al., 2012), as mentioned in Section 1.1. We describe the MIT data in more detail and motivate the model selected. We illustrate one of the advantages of our inferred network over the network snapshots provided by the survey data.

### 4.1 The Data

The data consists of Bluetooth proximity information collected from the phones of 57 MIT undergraduates living in a single dormitory over the period of a school year. Although political leanings and health measures were collected by the surveys, most basic personal information about the students is unavailable (e.g. sex), with the exception of what floor each student lived on and what year of study they were in.

After appropriate data cleaning, as detailed in Appendix A.1, the dataset consists of 66,432 Bluetooth-based proximity interactions among the 1,596 dyads. We model each dyad only over the interval when both students’ phones are active (sending/receiving Bluetooth signals) in the proximity logs. Summaries of the data are plotted in Figure 3 and reveal strong temporal dependencies. The left panel shows the rate of logged interactions increases throughout the school year. The center and right panels suggests strong daily and weekly periodicity: the majority of interactions occur at night, when students are more likely to be in their dormitory, and interactivity during the weekend appears to deviate from weekday behavior, with less of a disparity between daytime and nighttime activity levels.




**Figure 3:** Descriptive plots of the MIT Bluetooth proximity interactions. Left: Distribution of interactions over the school year. Center: Interaction frequencies within a week, aggregated from October to December for each dyad. Right: Interaction frequencies within a week, aggregated from April to May.

## 4.2 The Model

Figure 3 reveals both weekly periodicity patterns and an increase in activity level as the year progresses. To model the former, we approximate the weekly periodicity with sine waves, using the number of waves to trade off between simplicity and flexibility. We approximate the trend of increased overall interactivity over the school year by splitting up the year into three terms: first semester (Oct. - Dec.), second semester before spring break (Jan. - March), and second semester after spring break (April - May), and allowing for different activity levels in each period. The similarities between the center and right panels in Figure 3 suggest that the pattern of weekly periodicity is largely unrelated to the increase in overall interactivity over the school year.

Dyadic covariates and friendship relations between students are associated with weekly patterns of interactivity and thus are important to explore further in order to build a reasonable model. One unique aspect of the MIT data is that we can leverage the survey data in model construction. At the level of the dyad, only a weak relation exists between the number of interactions recorded and how often friendship was reported in surveys (see Figure 12 in Appendix A.1). However, when aggregating across dyads, we find significant differences in the interactivity patterns between reported friends and non-friends. Using survey data as a proxy for friendship, we compare the weekly patterns for reported friends/non-friends and individuals who live on the same/different floors, where we consider a dyad to be friends if at least one student named the other as a “close friend”. In the right panel of Figure 4, non-friends that live on the same floor tend to interact with higher propensity than non-friends that on different floors. However, while their overall interactivity levels differ, individuals



**Figure 4:** A comparison of weekly interactivity patterns by survey-reported friendship and whether or not pairs of individuals reside on the same floor. Daytime is highlighted with grey bars. Intensities are normalized so that non-friends on different floors have mean intensity of one. In order to avoid potential friendship dynamics, we consider a pair of students “friends” if at least one of the two identified themselves as “close friends” in four of the five surveys conducted. We consider a pair of students “non-friends” if they jointly identified themselves as “close friends” once or less.

in these two groups interact in similar weekly patterns. The left and center panels suggest that friends interact more than non-friends regardless of floor status, although the relative compounding effect is stronger for friends who live on different floors. Additionally, we find some evidence that friends may interact more than non-friends during the daytime (defined here to be 8 AM to 5 PM). Since the students all live in the same dormitory, encounters during the day, when more conscious effort is required to be in close proximity to others, may be especially indicative of friendship.

Using the notation in (2), we construct the following interaction model:

$$w(t, \text{floor}(i), \text{floor}(j)) = c_{0,\text{term}} \left( 1 + c_1 \mathbf{1}_{\text{floor}(i)=\text{floor}(j)} \right) \left( k_0 + \sum_{l=1}^3 k_{l,1} \cos(k_{l,2}t + k_{l,3}) \right) \quad (6)$$


$$m(\text{floor}(i), \text{floor}(j)) = c_{2,1} \mathbf{1}_{\text{floor}(i) \neq \text{floor}(j)} + c_{2,2} \mathbf{1}_{\text{floor}(i)=\text{floor}(j)} \quad (7)$$

$$a(t, \text{floor}(i), \text{floor}(j)) = c_{0,\text{term}} \left( 1 + c_1 \mathbf{1}_{\text{floor}(i)=\text{floor}(j)} \right) c_3 \mathbf{1}_{\text{daytime}}(t) \quad (8)$$

In the baseline  $w$ , three sets of sinusoidal terms model the weekly behavior while multipliers for school term and floor govern the overall levels of interactivity. We characterize friendship by an increase in daytime activity relative to the baseline pattern of weekly periodicity and an overarching multiplicative increase in intensity, dependent on floor status. Based on Figures 3 and 4, we assume this increased propensity to interact during the day for friends is subject to the same term- and floor-dependent modifiers that characterize general interactivity between dyads and use the same multipliers  $c_{0,\text{term}}$  and  $c_1$  in the daytime adjustment (8) that govern

the overall interactivity in the baseline (6).

We can similarly use the survey data to inform our model for the network paths, treating the reported relationships as observations of CTMCs at the survey dates. In Figure 5, living



**Figure 5:** Empirical CTMC parameter values for the four possible dyadic floor/year pairings derived from survey data. Sparsity  $s$  is estimated by the proportion of dyads who report friendship at each date, while transition parameter  $q$  is derived from the proportion of dyads which change reported friendship between consecutive survey dates.

on the same floor and being in the same year have independent, multiplicative effects on the empirical probability of friendship for dyads. Additionally, these proportions are constant throughout the school year and changes between survey dates are consistent with a single, low transition rate  $q$ . We model the network paths as CTMCs with sparsity

$$s(\text{floor}(i), \text{floor}(j), \text{year}(i), \text{year}(j)) = (1 + s_1 \mathbf{1}_{\text{floor}(i)=\text{floor}(j)}) (1 + s_2 \mathbf{1}_{\text{year}(i)=\text{year}(j)}) s_0 \quad (9)$$


and constant transition rate  $q$ . Introducing this structure into the sparsity model encourages block model behavior in the underlying network, where students who live on the same floor or are in the same year are more likely to be friends with one another and hence form “blocks” of friendship.

Before we can obtain estimates of the student network, we must assign priors to the parameters of our model. We observe a strong cyclical pattern in communications (associated with, for example, day and night), so we fix  $k_{i,2}$  using a Fourier transform on the observed weekly pattern of interactivity. We set the interactivity multiplier for the first school term equal to one ( $c_{0,t1} = 1$ ) to ensure the identifiability of these multipliers. We specify weakly informative priors on the remaining parameters based on support matching; additional details on these prior distributions are presented in Appendix A.2.

### 4.3 Network Analysis

We used Stan to generate 9,000 samples (after burn-in) of the parameters from the posterior distribution. Details of the sampling procedure are provided in Appendix A.2.

One of the key benefits of the inferred network is that it evolves in continuous-time, allowing for a more detailed examination of network evolution. For example, in the left two panels of Figure 6 we plot two snapshots of the MIT network as self-reported by students. In the first survey date shown, December 2008, student #2 (highlighted, towards the top




**Figure 6:** Snapshots of networks of the MIT students at two survey dates. Students are colored by the floor they lived on. For the reported network, students are considered friends if either reports the other as a “close friend.” In the case of the inferred networks, students are colored white if their phones are inactive and we can no longer infer relations for them. Uncertainty in the inferred relations is quantified by the opacity of the ties. In all snapshots, student #2 is highlighted, as well as his/her relations.

left) was reported to be friends with three other individuals. By March of the following year, there were eight friendships involving this student. Between the two dates, six friendships were formed and one dissolved. Yet, with only this information, we have little insight into how these changes came about and cannot test theories about the social processes underlying these changes, which are of particular interest for actor-oriented models (Snijders, 1996).

The continuous-time evolution of our inferred network provides greater insight into potential social theories. In the right panels of Figure 6 we plot snapshots of the inferred network at the same two survey dates. We observe similar behavior with student #2, albeit with different individuals; we inferred three friendships in December and seven in March, corresponding to five formed friendships and one dissolved. Since the inferred network evolves in continuous-time, we can examine snapshots of the network at any time in between the survey dates to get a better sense of how the network evolves and see how the interactions influence the formation/dissolution of inferred ties.

In Figure 7, we plot a subset of the inferred network for six additional dates between the survey dates. An animated series of all daily snapshots between the survey dates can be found at <https://git.io/viTtS>. These intermediate snapshots of the network provide evidence of triadic closure in this community, as friendships do not tend to form randomly



**Figure 7:** Snapshots of the inferred network between December 13th, 2008 and March 5th, 2009 for a selected subset of students. Students are colored by the floor they lived on. Edge widths for each dyad denote the estimated mean probability of friendship at the given time.

but rather appear with higher propensity between actors who have mutual connections. We would have missed this behavior without the ability to examine the evolution of the network at finer resolutions; only once we were able to examine intermediate evolutions of the network were we able to find evidence of this phenomenon.

## 5 Proximity Interactions among Barn Swallows

We now turn to the behavior of North American barn swallows (*Hirundo rustica erythrogaster*). Unlike in the MIT student data, there is no survey data to guide our selection of the functions  $w$ ,  $m$ ,  $a$ ,  $s$ , and  $q$  underpinning the model.

### 5.1 The Data


Telemetry data were gathered for 17 barn swallows monitored from 06:00-09:00 and 17:00-20:00 (periods of high barn swallow activity) over three consecutive days during mating season (Levin et al., 2015, 2016). Barn swallows were outfitted with tracking devices which logged the presence of nearby swallows every 20 seconds. After data cleaning, detailed in Appendix B.1, the resulting dataset has 1009 interactions among its 136 dyads. Of the barn swallows studied, 10 were male and 7 were female. Various other physical characteristics of the birds are also available, including mass, tail streamer length, and plumage color.



## 5.2 The Model

We choose to omit various observable covariates known to be associated with barn swallow social behavior from our model. For example, our model does not explicitly encode for factors that influence mate selection, such as the female preference to mate with males with darker plumage color (Safran and McGraw, 2004; Safran et al., 2005). Instead, our strategy is to develop a parsimonious model of the marginal network behavior without considering observables endogenous to network formation. This strategy, which may be necessary for less well-studied communities, provides an overall characterization of barn swallow social behavior that is useful in its own right and, as we show in the results, can also be correlated with observable features after model fitting.

In Figure 8, we see that both sex and time play important roles in understanding interaction behavior between swallows. The left panel shows that female birds interact very little with one another, yet interact quite often with males. In contrast, male birds regularly interact with each other. We take these differences in activity levels to represent behavioral discrepancies in the types of relationships formed by the various sex pairings rather than differences in the propensity for forming these relationships themselves. The right panel shows the rates of interaction during the periods of observation are non-constant. Differences in these rates may be due to factors outside of our interest, such as the weather and other variable environmental stimuli, which are unlikely to impact network dynamics.



**Figure 8:** Descriptive statistics of the interaction data. Left: Histogram of the number of interactions for each dyad, color coded by the sex pairing of each dyad. Right: Histogram of the number of interactions over time, aggregated across dyads. Interactions are also color coded by its dyads' sex pairing.

Under the assumption that swallows without social ties interact at a rate independent of their sex pairing, we choose to model the baseline activity rate for any pair at time  $t$  as

$$w(t) = k_1 \mathbf{1}_{7/19 \text{ PM}} + k_2 \mathbf{1}_{7/20 \text{ AM}} + \dots + k_8 \mathbf{1}_{7/23 \text{ AM}}. \quad (10)$$



In this model, baseline activity rate is constant within each observational period but varies from period to period. This allows for flexibility in the temporal trend without imposing a strict parametric trend.

In contrast, we believe that barn swallows have different types of social relationships depending on the sexes of the birds involved and thus sex pairing will affect the activity levels of familiar birds. We embed this phenomenon by allowing for different multiplicative increases in interaction rates depending on sex pairing. Omitting an additive effect  $a$  to the baseline intensity, we assume that birds with social ties interact with the same temporal patterns as birds without ties. The intensity rate of the  $(i, j)$ th bird pair at time  $t$  is then given by


$$\lambda_{ij}(t) = (1 + (c_{FF}\mathbf{1}_{FF} + c_{MF}\mathbf{1}_{MF} + c_{MM}\mathbf{1}_{MM})y_{ij}(t))w(t), \quad (11)$$

where  $y_{ij}(t)$  is an indicator for whether or not  $(i, j)$  are socially connected at time  $t$ .

In accordance with our strategy to not explicitly encode observables into the formulation of the network, we model each dyad's connection  $y_{ij}(t)$  with i.i.d. CTMCs with common sparsity  $s$  and common transition parameter  $q$ . We complete the model by assigning weakly informative priors based on support matching to the parameters  $\{k_1, \dots, k_8, c_{FF}, c_{MF}, c_{MM}, s, q\}$ . Detailed priors are given in Appendix B.1.

### 5.3 Network Analysis

We considered 9,000 samples (after burn-in) obtained from Stan (see Appendix B.1 for computational details). A full animation of our inferred network, alongside the raw encounter network, is available at <https://git.io/viTtP>.




**Figure 9:** Estimated mean paths grouped by sex pairing.

In Figure 9, we group the estimated mean network paths by sex pairing. In the left panel, we see the estimated social connections between female birds are generally static, without

large changes in the probability of social ties over the course of the study. Unfortunately, some of the tracking devices on the female swallows became inactive early in the study, obscuring potential changes in the latter half of the study.

Displayed in the right panel of Figure 9, pairs of male swallows exhibited a very diverse set of behaviors. In stark contrast to pairs of female swallows, many male pairs had dynamic associations in which their interactivity changed dramatically over the time of the study. The probability of a social tie fluctuated by at least 50% over the three day period for slightly more than 20% of possible male pairs. The differences in behavior may be attributable to the aggression of males during mating season, when they are particularly prone to engaging in territorial conflict with one another. At the same time, other pairs had relatively stable social ties or lack thereof. For example, we estimated 10 out of the possible 45 pairs to be socially connected with probability 0.8 or higher when averaged over the course the study. These relations are plotted in Figure 10. We observe a clique of four male swallows, as well




**Figure 10:** *Network of male swallows. Nodes are labeled with each bird’s tag identification number. Edges denote “strong” relationships, between birds who are predicted to be socially connected with probability 0.8 or higher on average over the study.*

as four other males with fewer relations to the clique. The stable behavior between these males may be evidence of extended territorial conflict or more complex social structure, such as unmated males “attending” nests of other birds in the colony (Crook and Shields, 1987).

Both strong stability and some dynamics characterize the behavior of male/female pairs, displayed in the middle panel of Figures 9. We consistently classify about 75% of the possible male/female pairings to have no relation. The remaining 25% displayed a wide range of relational probabilities and dynamics. Among the latter were four disjoint pairs which averaged probability of relationships of 0.8 or higher. While barn swallows are sexually polygamous, they are socially monogamous and form pair-bonds for the breeding season; the behavior of these four pairs suggest that they may be such pairs.

Previous studies of European barn swallows established that females prefer males with longer tail streamers (Moller et al., 1998), while in the United States it has been documented that females prefer males with darker ventral plumage (Safran et al., 2005). Despite not encoding sexual selection into our model, we find evidence congruent with both theories for the studied North American barn swallows. Under the assumption that male/female relationships signal mating preferences, the males which were most successful in establishing bonds with females tended to have longer tail streamers and darker ventral plumage. Comparing the nine males for which tail streamer length was recorded, the three males with highest probability of relations with females had the three longest tails (among the nine). This finding, illustrated in Figure 11, is quite robust to the choice of cutoff dividing males by tail streamer length. A similar relationship held to a lesser extent for ventral coloring, which was strongly correlated ( $r=0.79$ ) with tail streamer length for the studied males.



**Figure 11:** Male/female relationships among barn swallows. Red nodes denote females, while blue and black nodes denote males. Edges are weighted by average probability of a social relation, and estimated probabilities lower than 0.1 are omitted. In the left graph, males are shaded by tail streamer length, with darker color corresponding to longer tails (except for black, which corresponds to a missing value). In the right graph, males are shaded according to the color of their ventral plumage, with darker blue corresponding to darker coloration.

## 6 Discussion

The proliferation of continuous-time relational event data presents both new opportunities and challenges for social network research. Event data represent observed behavior, which fundamentally differs from reported behavior, the traditional domain of survey data (Killworth and Bernard, 1976; Eagle et al., 2009). Many of the network analysis tools developed for survey data are not appropriate for relational event data (Howison et al., 2011), and new

methods must be proposed. Simultaneously, a better understanding of the data reliability issues surrounding the data collection process is required. Most importantly, these data expose numerous ethical issues related to how the data should be used and how to ensure privacy of individuals.

We have presented a hierarchical network model in which relational events are modeled with Poisson processes governed by dyadic covariates and a dynamic latent network. Our model acknowledges the intrinsic nature of relational event data as continuous-time interactions only indirectly indicative of underlying dynamic network relations.

Depending on the intended application, further developments of the model may allow for more precise estimates of the latent network or more fully capturing dependence between events. For example, in the MIT study considered in Section 4, calls and text messages between participants were also logged. While these sets of data are less rich than the proximity data (there are only seven thousand phone calls and thirteen hundred text messages, as opposed to sixty thousand proximity interactions), integrating multiple data sources through the use of multiple connected point processes for each dyad may yield more accurate network estimates. Alternatively, multiple types of relationships may exist between the students, which we may wish to index with multiple latent networks. As another example, when modeling human telecommunication data where reciprocity between interactions play a larger role in explaining dependencies between events (e.g. text messaging and email), self-exciting Hawkes processes may be more suitable than the Poisson processes. The structure of relational event data, its relationships to an underlying network, and the meaning of the network will vary considerably across applications. Hence, domain knowledge should be used whenever possible to customize the model to target capturing relevant and interesting structure.

## A Appendix: MIT Social Evolution Data

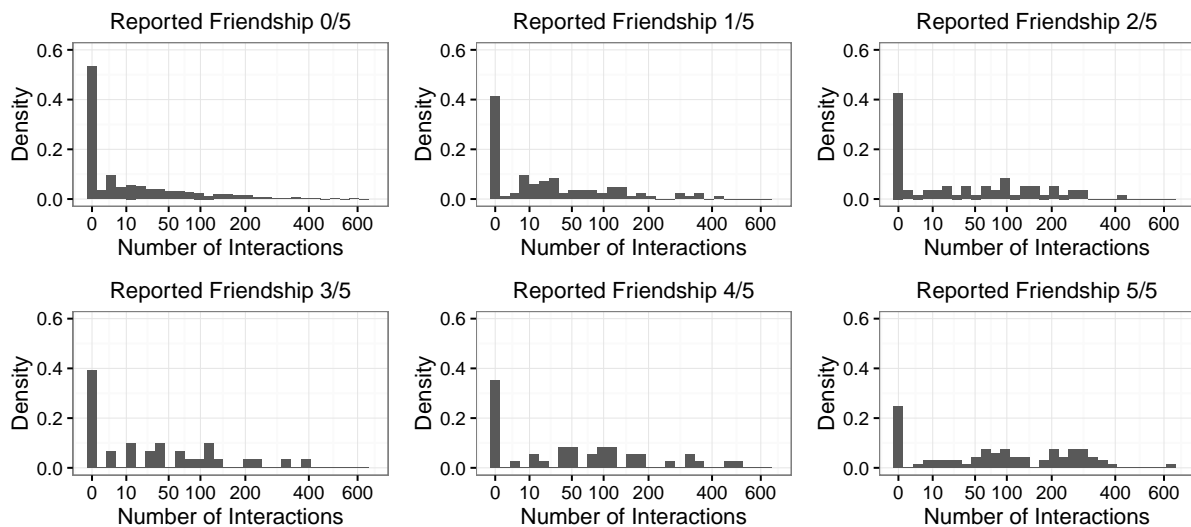
### A.1 Data Considerations

Participants in the study were given smartphones outfitted with an application that used Bluetooth to scan for other nearby (within 10 meters) smartphones approximately every six minutes. However, the application only logged the presence of other devices set as “discoverable.” Due to this and other technical issues, many of the reported encounters were non-reciprocal and it is possible many encounters went unreported. In addition, a fraction of the logged encounters occurred between phones on different floors due to vertical

proximity. These encounters are not of interest, and are identifiable as being estimated by the application to have low probability of being on the same floor using WLAN location data. We chose to drop all encounters with probability 0.2 or lower. In addition, we only considered interactions that occur during the Fall (September 3rd to December 13th) and Spring sessions (January 5th to March 20th and March 30th to May 15th), notably excluding winter and spring breaks, which were extended periods when school is not in session. Please refer to [Madan et al. \(2012\)](#) for further details about the data.

We converted the proximity logs into undirected (reciprocal) interaction data by combining the directed proximity logs for pairs of individuals, taking encounters logged in either phone. Logs that were time-stamped within 30 minutes of one another were taken to correspond to the same interaction.

Defining a dyad to have reported as friends in a survey if either marks the other as a ‘close friend,’ we observed significant dissonance between the survey data and interaction data on a dyadic level. Over twenty percent of dyads who reported friendship in all five surveys did not have any recorded interactions, compared to just under fifty percent for dyads who never reported friendship. The weak relation between the sources of data was robust to different choices of defining friendship in the survey data and raises validity concerns about both sets of data.



**Figure 12:** Distributions of the total number of recorded interactions per dyad, where the data is partitioned by the fraction of times each dyad reported friendship in the five surveys. We consider a dyad to have reported as friends in a survey if either individuals marks the other as a ‘close friend.’

## A.2 Priors and Model Fitting

We assign the following independent priors to the parameters of the model specified by (6)-(9): (where applicable, units are in hours)

$$\begin{aligned}
\{c_{0,t2}, c_{0,t3}\} &\sim \text{Exp}(1) & c_1 &\sim \text{Exp}(1) \\
\{c_{2,1}, c_{2,2}\} &\sim \text{Exp}(100) & c_3 &\sim \text{Exp}(1) \\
k_0 &\sim \text{Exp}(1) & \{k_{i,1} : i \in \{1, 2, 3\}\} &\sim \text{Exp}(1) \\
\{k_{i,3} : i \in \{1, 2, 3\}\} &\sim \text{Uniform}(0, 2\pi) & s_0 &\sim \text{Beta}(1, 49) \\
s_1 &\sim \text{Exp}(1) & s_2 &\sim \text{Exp}(1) \\
q &\sim \text{Exp}(1e10)
\end{aligned}$$


We used Stan to generate 10,000 samples using a single MCMC chain, initializing each of the parameters at its prior mean and discarding the first 1,000 as burn-in. This sampling procedure took approximately 24 hours on a standard laptop. We assessed convergence by visually examining the trace plots shown in Figure 13.

## B Appendix: Barn Swallow Data

### B.1 Data Considerations

Every 20 seconds, the proximity loggers attached to each barn swallow logged the presence of other barn swallows within 5 meters. This threshold was chosen by the researchers since it represented the average distance between nests at the observation site. The devices split logged encounters between birds that lasted more than 5 minutes into 5 minute intervals. Unfortunately, the proximity loggers were imperfect; sometimes encounters recorded by one bird were not recorded by the other and the batteries of some devices died (starting on July 22nd) before the end of the observation period. See [Levin et al. \(2015\)](#) for a more comprehensive description of the data.

We accounted for asymmetries and extended encounters in the encounter log. To ensure encounters that last longer than 5 minutes were treated as a single interaction, encounters separated by thirty seconds or less were merged. To compensate for the imperfect behavior of the tracking devices, we combined the logged encounter history for every pair of birds, taking encounters logged in either birds' tracker as an interaction and merging overlapping encounters. Lastly, to help avoid issues related to battery-life, we only modelled each dyad over the interval when the trackers of both birds recorded activity.



**Figure 13:** Trace plots for the MIT model parameters. Values obtained during the burn-in period are discarded.

## B.2 Priors and Model Fitting

We assign the following priors to the parameters of our model: (units for  $k_i$  and  $q$  are provided in hours)

$$\begin{aligned} \{c_{FF}, c_{MF}, c_{MM}\} &\sim \text{Exp}(2) & \{k_i : i \in \{1, \dots, 8\}\} &\sim \text{Exp}(1) \\ s &\sim \text{Beta}(1, 9) & q &\sim \text{Exp}(1000) \end{aligned}$$

Using Stan, we ran a single MCMC chain for 10,000 iterations and discarded the first 1,000 as burn-in. The parameters were initialized at the means of their respective prior distributions, and the entire sampling procedure took approximately 6 minutes. As with the previous application, we assessed convergence by examining trace plots, provided in Figure 14.

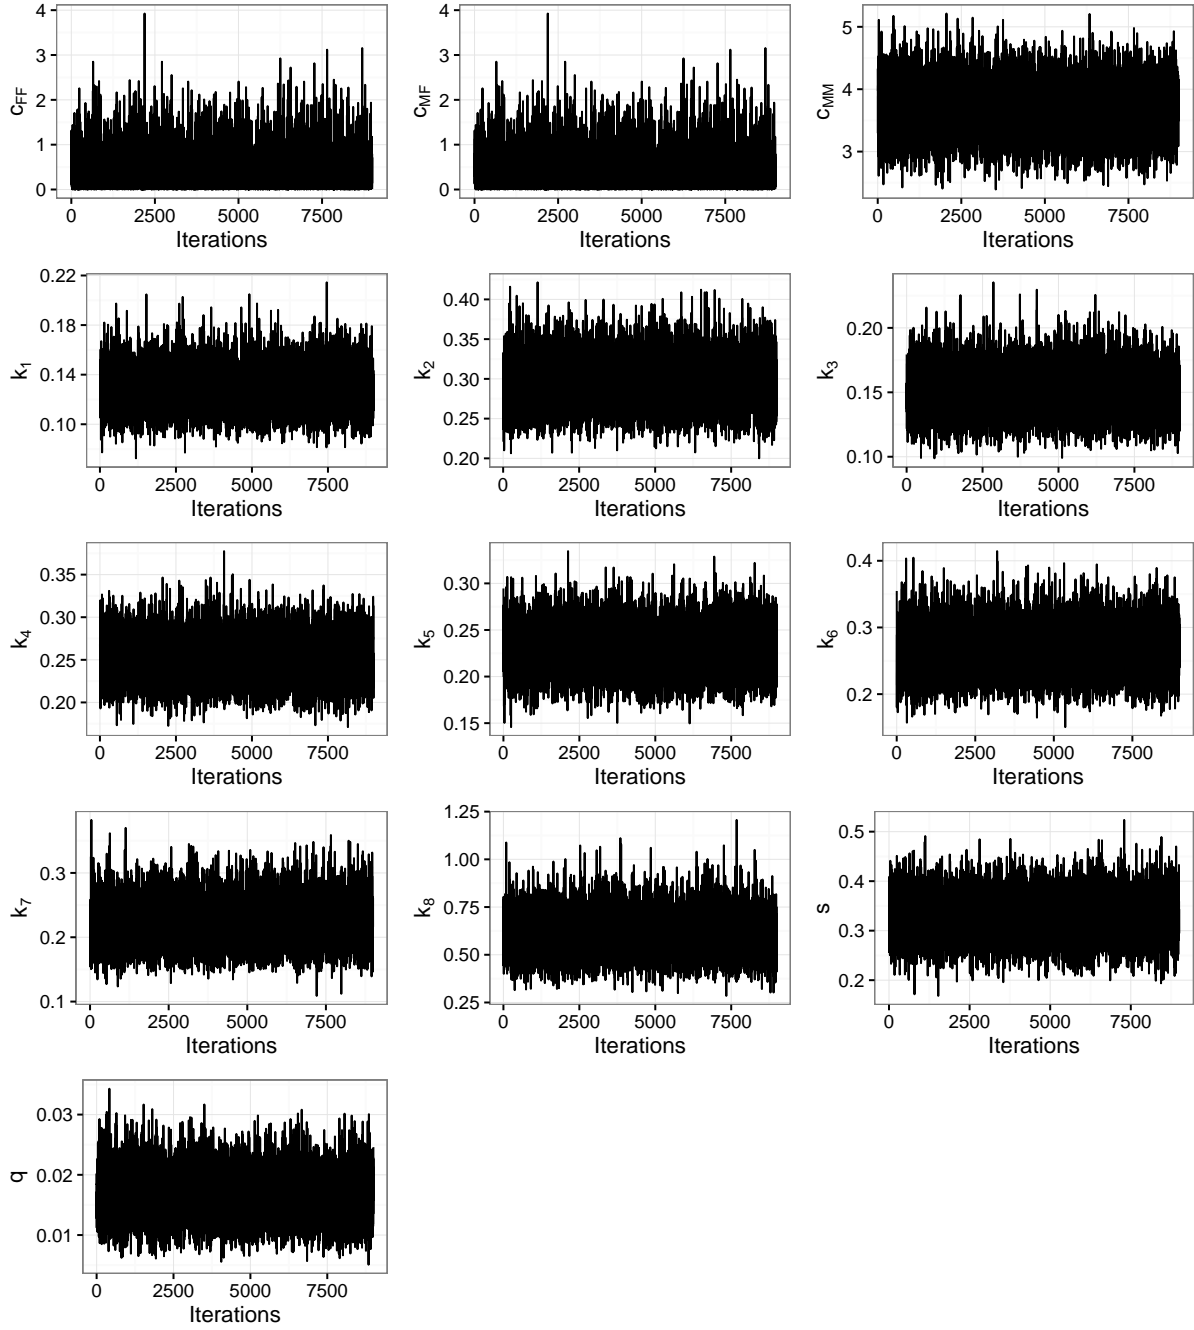
## Acknowledgements

The authors would like to thank Iris Levin and Henry Scharf for helpful comments and suggestions on an earlier version of the manuscript, and Iris Levin and Rebecca Safran for providing the barn swallow interaction data. This work was partially supported by the National Science Foundation under Grant Number SES-1461495 to Fosdick and Grant Number SES-1559778 to McCormick. This material is based upon work supported by, or in part by, the U. S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number W911NF-12-1-0379.

## References

- Butts CT. A relational event framework for social action. *Sociological Methodology* 2008; **38** (1):155–200. DOI: 10.1111/j.1467-9531.2008.00203.x.
- Blumenstock J. Inferring patterns of internal migration from mobile phone call records: evidence from Rwanda. *Information Technology for Development* 2012; **18**(2):107–125. DOI: 10.1080/02681102.2011.643209.
- Sadilek A, Kautz H, Silenzio V. Modeling spread of disease from social interactions. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media* 2012; 322–329.





**Figure 14:** Trace plots for barn swallow model parameters. Values obtained during the burn-in period are discarded.

- Rutz C, Hays GC. New frontiers in biologging science. *Biology Letters* 2009; **5**(3):289–292. DOI: 10.1098/rsbl.2009.0089.
- Watts DJ. A twenty-first century science. *Nature* 2007; **445**(7127):489–489. DOI: 10.1038/445489a.
- Sarkar P, Moore AW. Dynamic social network analysis using latent space models. *SIGKDD Explorations: Special Edition on Link Mining* 2005; **7**(2):31–40. DOI: 10.1145/1117454.1117459.
- Durante D, Dunson DB. Nonparametric Bayes dynamic modeling of relational data. *Biometrika* 2014; **101**(4):883–898. DOI: 10.1093/biomet/asu040.
- Krivitsky PN, Handcock MS. A separable model for dynamic networks. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 2014; **76**(1):29–46. DOI: 10.1111/rssb.12014.
- Sewell DK, Chen Y. Latent space models for dynamic networks. *Journal of the American Statistical Association* 2015; **110**(512):1646–1657. DOI: 10.1080/01621459.2014.988214.
- Snijders T. Models for longitudinal network data. *Models and Methods in Social Network Analysis* 2005; **11**:215–247.
- Sewell DK, Chen Y. Latent space models for dynamic networks with weighted edges. *Social Networks* 2016; **44**:105–116. DOI: 10.1016/j.socnet.2015.07.005.
- Sulo R, Berger-Wolf T, Grossman R. Meaningful selection of temporal resolution for dynamic networks. In *Proceedings of the Eighth Workshop on Mining and Learning with Graphs* 2010; 127–136. DOI: 10.1145/1830252.1830269.
- Sekara V, Stopczynski A, Lehmann S. Fundamental structures of dynamic social networks. In *Proceedings of the National Academy of Sciences of the United States of America* 2016; Forthcoming. DOI: 10.1073/pnas.1602803113.
- Simma A, Jordan MI. Modeling events with cascades of Poisson processes. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence* 2010; 546–555.
- Blundell C, Heller KA, Beck JM. Modelling reciprocating relationships with Hawkes processes. In *Advances in Neural Information Processing Systems 25* 2012; 2600–2608.

- Perry PO, Wolfe PJ. Point process modelling for directed interaction networks. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 2013; **75**(5):821–849. DOI: 10.1111/rssb.12013.
- Hawkes AG. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 1971; **58**(1):83–90. DOI: 10.1093/biomet/58.1.83.
- Linderman SW, Adams RP. Discovering latent network structure in point process data. In *Proceedings of the 31st International Conference on Machine Learning* 2014; 1413–1421.
- DuBois C, Smyth P. Modeling relational events via latent classes. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 2010; 803–812. DOI: 10.1145/1835804.1835906.
- Scharf HR, Hooten MB, Fosdick BK, Johnson DS, London JM, Durban JW. Dynamic social networks based on movement. *ArXiv e-prints* 2015; 1–27.
- Madan A, Cebrian M, Moturu ST, Farrahi K, Pentland AS. Sensing the health state of a community. *IEEE Pervasive Computing* 2012; **11**(4):36–45. DOI: 10.1109/MPRV.2011.79.
- Dong W, Pentland AS, Heller KA. Graph-coupled HMMs for modeling the spread of infection. *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence* 2012; 227–236.
- Madan A, Moturu ST, Lazer D, Pentland AS. Social sensing: Obesity, unhealthy eating and exercise in face-to-face networks. In *Wireless Health 2010* 2010; 104–110.
- Levin II, Zonana DM, Burt JM, Safran RJ. Performance of encounter tags: Field tests of miniaturized proximity loggers for use on small birds. *PLoS ONE* 2015; **10**(9). DOI: 10.1371/journal.pone.0137242.
- Levin II, Zonana DM, Fosdick BK, Song SJ, Knight R, Safran RJ. Stress response, gut microbial diversity and sexual signals correlate with social interactions. *Biology Letters* 2016; **12**(6):20160352.
- Rabiner LR. A tutorial on hidden Markov models and selected applications in speech recognition. In *Proceedings of the IEEE* 1989; **77**(2):257–286. DOI: 10.1109/5.18626.
- Carpenter B, Gelman A, Hoffman M, Lee D, Goodrich B, Betancourt M, Brubaker MA, Li P, Riddell A. Stan : A probabilistic programming language. *Journal of Statistical Software* 2016; **VV**(II).

- Snijders T. Stochastic actor-oriented models for network change. *The Journal of Mathematical Sociology* 1996; **21**(1-2):149–172. DOI: 10.1080/0022250X.1996.9990178.
- Safran RJ, McGraw KJ. Plumage coloration, not length or symmetry of tail-streamers, is a sexually selected trait in North American barn swallows. *Behavioral Ecology* 2004; **15**(3): 455–461. DOI: 10.1093/beheco/arh035.
- Safran RJ, Neuman CR, McGraw KJ, Lovette IJ. Dynamic paternity allocation as a function of male plumage color in barn swallows. *Science* 2005; **309**(5744):2210–2. DOI: 10.1126/science.1115090.
- Crook JR, Shields WM. Non-parental nest attendance in the barn swallow (*Hirundo rustica*): helping or harassment? *Animal Behaviour* 1987; **35**(4):991–1001. DOI: 10.1016/S0003-3472(87)80156-2.
- Moller AP, Barbosa A, Cuervo JJ, de Lope F, Merino S, Saino N. Sexual selection and tail streamers in the barn swallow. *Proceedings of the Royal Society B: Biological Sciences* 1998; **265**(1394):409–414. DOI: 10.1098/rspb.1998.0309.
- Killworth PD, Bernard HR. Informant accuracy in social network data. *Human Organization* 1976; **35**(3):269–286.
- Eagle N, Pentland AS, Lazer D. Inferring friendship network structure by using mobile phone data. In *Proceedings of the National Academy of Sciences of the United States of America* 2009; **106**(36):15274–15278. DOI: 10.1073/pnas.0900282106.
- Howison J, Wiggins A, Crowston K. Validity issues in the use of social network analysis with digital trace data. *Journal of the Association for Information Systems* 2011; **12**(12): 767–797.